

"Express Mail" Mailing Label No. EV352470158US

PATENT APPLICATION

5

10 SELECTIVELY MONITORING STORES TO SUPPORT TRANSACTIONAL PROGRAM EXECUTION

Inventors: Marc Tremblay, Quinn A. Jacobson and Shailender Chaudhry

Related Application

[0001] This application hereby claims priority under 35 U.S.C. §119 to
20 U.S. Provisional Patent Application No. 60/447,128, filed on 13 February 2003,
entitled “Transactional Memory,” by inventors Shailender Chaudhry, Marc
Tremblay and Quinn Jacobson (Attorney Docket No. SUN-P9322PSP).

[0002] The subject matter of this application is additionally related to the subject matter in a co-pending non-provisional U.S. patent application by the same inventors as the instant application and filed on the same day as the instant application entitled, “Selectively Monitoring Loads to Support Transactional Program Execution,” having serial number TO BE ASSIGNED, and filing date TO BE ASSIGNED (Attorney Docket No. SUN-P9327-MEG).

BACKGROUND

Field of the Invention

[0003] The present invention relates to techniques for improving the performance of computer systems. More specifically, the present invention relates to a method and an apparatus for selectively monitoring loads to support transactional program execution.

Related Art

10 [0004] Computer system designers are presently developing mechanisms to support multi-threading within the latest generation of Chip-Multiprocessors (CMPs) as well as more traditional Shared Memory Multiprocessors (SMPs). With proper hardware support, multi-threading can dramatically increase the performance of numerous applications. However, as microprocessor performance 15 continues to increase, the time spent synchronizing between threads (processes) is becoming a large fraction of overall execution time. In fact, as multi-threaded applications begin to use even more threads, this synchronization overhead becomes the dominant factor in limiting application performance.

[0005] From a programmer's perspective, synchronization is generally 20 accomplished through the use locks. A lock is typically acquired before a thread enters a critical section of code, and is released after the thread exits the critical section. If another thread wants to enter the same critical section, it must acquire the same lock. If it is unable to acquire the lock, because a preceding thread has grabbed the lock, the thread must wait until the preceding thread releases the lock. 25 (Note that a lock can be implemented in a number of ways, such as through atomic operations or semaphores.)

[0006] Unfortunately, the process of acquiring a lock and the process of releasing a lock are very time-consuming in modern microprocessors. They involve atomic operations, which typically flush the load buffer and store buffer, and can consequently require hundreds, if not thousands, of processor cycles to 5 complete.

[0007] Moreover, as multi-threaded applications use more threads, more locks are required. For example, if multiple threads need to access a shared data structure, it is impractical for performance reasons to use a single lock for the entire data structure. Instead, it is preferable to use multiple fine-grained locks to 10 lock small portions of the data structure. This allows multiple threads to operate on different portions of the data structure in parallel. However, it also requires a single thread to acquire and release multiple locks in order to access different portions of the data structure.

[0008] In some cases, locks are used when they are not required. For 15 example, many applications make use of “thread-safe” library routines that use locks to ensure that they are “thread-safe” for multi-threaded applications. Unfortunately, the overhead involved in acquiring and releasing these locks is still incurred, even when the thread-safe library routines are called by a single-threaded application.

[0009] Applications typically use locks to ensure mutual exclusion within 20 critical sections of code. However, in many cases threads will not interfere with each other, even if they are allowed to execute a critical section simultaneously. In these cases, mutual exclusion is used to prevent the unlikely case in which threads actually interfere with each other. Consequently, in these cases, the 25 overhead involved in acquiring and releasing locks is largely wasted.

[0010] Hence, what is needed is a method and an apparatus that reduces the overhead involved in manipulating locks when accessing critical sections of code.

[0011] One technique to reduce the overhead involved in manipulating locks is to “transactionally” execute a critical section, wherein changes made during the transactional execution are not committed to the architectural state of the processor until the transactional execution successfully completes. This technique is described in related U.S. Patent Application No. 10/439,911, entitled, “Method and Apparatus for Avoiding Locks by Speculatively Executing Critical Sections,” by inventors Shailender Chaudhry Marc Tremblay and Quinn A. Jacobson, filed on 16 May 2003 (Attorney Docket No. SUN-P9322-MEG).

[0012] During transactional execution, load and store operations are modified so that they mark cache lines that are accessed during the transactional execution. This allows the computer system to determine if an interfering data access occurs during the transactional execution, in which case the transactional execution fails, and results of the transactional execution are not committed to the architectural state of the processor.

[0013] Unfortunately, problems can arise while marking cache lines. If a large number of lines are marked, false failures are likely to occur when accesses that appear to interfere with each other do not actually touch the same data items in a cache line. Furthermore, the marked cache lines cannot be easily moved out of cache until the transactional execution completes, which also causes performance problems.

[0014] Also, since store operations need to be buffered during transactional execution, transactional execution will sometimes be limited by the number of available store buffers on the processor.

[0015] Hence, what is needed is a method and an apparatus that reduces the number of cache lines that need to be marked during transactional program execution.

5

SUMMARY

[0016] One embodiment of the present invention provides a system that selectively monitors load instructions to support transactional execution of a process, wherein changes made during the transactional execution are not committed to the architectural state of a processor until the transactional execution 10 successfully completes. Upon encountering a load instruction during transactional execution of a block of instructions, the system determines whether the load instruction is a monitored load instruction or an unmonitored load instruction. If the load instruction is a monitored load instruction, the system performs the load operation, and load-marks a cache line associated with the load instruction to 15 facilitate subsequent detection of an interfering data access to the cache line from another process. If the load instruction is an unmonitored load instruction, the system performs the load operation without load-marking the cache line.

[0017] In a variation on this embodiment, prior to executing the program, the system generates instructions for the program. During this process, the system 20 determines whether load operations that take place during transactional execution need to be monitored. The system then generates monitored load instructions for load operations that need to be monitored, and generates unmonitored load instructions for load operations that do not need to be monitored.

[0018] In a variation on this embodiment, the system determines whether a 25 load operation needs to be monitored by determining whether the load operation is directed to a heap, wherein loads from the heap need to be monitored and loads from outside the heap do not need to be monitored.

[0019] In a variation on this embodiment, the system determines whether a load operation needs to be monitored by examining a data structure associated with the load operation to determine whether the data structure is a “protected” data structure for which loads need to be monitored, or an “unprotected” data structure for which loads do not need to be monitored.

5 [0020] In a variation on this embodiment, the system determines whether a load operation needs to be monitored by allowing a programmer to determine if the load operation needs to be monitored.

10 [0021] In a variation on this embodiment, the system determines whether a load operation needs to be monitored by examining an op code of the load instruction.

15 [0022] In a variation on this embodiment, the system determines whether a load operation needs to be monitored by examining an address associated with the load instruction to determine whether the address falls within a range of addresses for which loads are monitored. Examining the address can involve comparing the address with one or more boundary registers. It can also involve examining a Translation Lookaside Buffer (TLB) entry associated with the address.

20 [0023] In a variation on this embodiment, if an interfering data access from another process is encountered during transactional execution of the block of instructions, the system discards changes made during the transactional execution and attempts to re-execute the block of instructions.

25 [0024] In a variation on this embodiment, if transactional execution of the block of instructions completes without encountering an interfering data access from another process, the system commits changes made during the transactional execution to the architectural state of the processor, and resumes normal non-transactional execution of the program past the block of instructions.

[0025] In a variation on this embodiment, an interfering data access can include: a store by another process to a cache line that has been load-marked by the process; and a load or a store by another process to a cache line that has been store-marked by the process.

- 5 **[0026]** In a variation on this embodiment, the cache line is load-marked in level 1 (L1) cache.

[0027] One embodiment of the present invention provides a system that selectively monitors store instructions to support transactional execution of a process, wherein changes made during the transactional execution are not
10 committed to the architectural state of a processor until the transactional execution successfully completes. Upon encountering a store instruction during transactional execution of a block of instructions, the system determines whether the store instruction is a monitored store instruction or an unmonitored store instruction. If the store instruction is a monitored store instruction, the system
15 performs the store operation, and store-marks a cache line associated with the store instruction to facilitate subsequent detection of an interfering data access to the cache line from another process. If the store instruction is an unmonitored store instruction, the system performs the store operation without store-marking the cache line.

- 20 **[0028]** In a variation on this embodiment, prior to executing the program, the system generates instructions for the program. During this process, the system determines whether store operations that take place during transactional execution need to be monitored. The system then generates monitored store instructions for store operations that need to be monitored, and generates unmonitored store
25 instructions for store operations that do not need to be monitored.

[0029] In a variation on this embodiment, the system determines whether a store operation needs to be monitored by determining whether the store operation

is directed to a heap, wherein stores from the heap need to be monitored and stores from outside the heap do not need to be monitored.

- [0030] In a variation on this embodiment, the system determines whether a store operation needs to be monitored by examining a data structure associated
5 with the store operation to determine whether the data structure is a “protected” data structure for which stores need to be monitored, or an “unprotected” data structure for which stores do not need to be monitored.

- [0031] In a variation on this embodiment, the system determines whether a store operation needs to be monitored by allowing a programmer to determine if
10 the store operation needs to be monitored.

[0032] In a variation on this embodiment, the system determines whether the store instruction is a monitored store instruction by examining an op code of the store instruction.

- [0033] In a variation on this embodiment, the system determines whether
15 the store instruction is a monitored store instruction by examining an address associated with the store instruction to determine whether the address falls within a range of addresses for which stores are monitored.

[0034] In a variation on this embodiment, the cache line is store-marked in the cache level closest to the processor where cache lines are coherent.

- [0035] In a variation on this embodiment, a store-marked cache line can indicate that: loads from other processes to the cache line should be monitored; stores from other processes to the cache line should be monitored; and stores to the cache line should be buffered until the transactional execution completes.
20

25

BRIEF DESCRIPTION OF THE FIGURES

[0036] FIG. 1 illustrates a computer system in accordance with an embodiment of the present invention.

[0037] FIG. 2 illustrates how a critical section is executed in accordance with an embodiment of the present invention.

[0038] FIG. 3 presents a flow chart illustrating the transactional execution process in accordance with an embodiment of the present invention.

5 [0039] FIG. 4 presents a flow chart illustrating a start transactional execution (STE) operation in accordance with an embodiment of the present invention.

10 [0040] FIG. 5 presents a flow chart illustrating how load-marking is performed during transactional execution in accordance with an embodiment of the present invention.

[0041] FIG. 6 presents a flow chart illustrating how store-marking is performed during transactional execution in accordance with an embodiment of the present invention.

15 [0042] FIG. 7 presents a flow chart illustrating how a commit operation is performed in accordance with an embodiment of the present invention.

[0043] FIG. 8 presents a flow chart illustrating how changes are discarded after transactional execution completes unsuccessfully in accordance with an embodiment of the present invention.

20 [0044] FIG. 9A presents a flow chart illustrating how monitored and unmonitored load instructions are generated in accordance with an embodiment of the present invention.

[0045] FIG. 9B presents a flow chart illustrating how monitored and unmonitored load instructions are executed in accordance with an embodiment of the present invention.

25 [0046] FIG. 10A presents a flow chart illustrating how monitored and unmonitored store instructions are generated in accordance with an embodiment of the present invention.

[0047] FIG. 10B presents a flow chart illustrating how monitored and unmonitored store instructions are executed in accordance with an embodiment of the present invention.

5

DETAILED DESCRIPTION

[0048] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general 10 principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

[0049] The data structures and code described in this detailed description 15 are typically stored on a computer readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital versatile discs 20 or digital video discs), and computer instruction signals embodied in a transmission medium (with or without a carrier wave upon which the signals are modulated). For example, the transmission medium may include a communications network, such as the Internet.

25 **Computer System**

[0050] FIG. 1 illustrates a computer system 100 in accordance with an embodiment of the present invention. Computer system 100 can generally include

any type of computer system, including, but not limited to, a computer system based on a microprocessor, a mainframe computer, a digital signal processor, a portable computing device, a personal organizer, a device controller, and a computational engine within an appliance. As is illustrated in FIG. 1, computer
5 system 100 includes processors 101 and level 2 (L2) cache 120, which is coupled to main memory (not shown). Processor 102 is similar in structure to processor 101, so only processor 101 is described below.

[0051] Processor 101 has two register files 103 and 104, one of which is an “active register file” and the other of which is a backup “shadow register file.”
10 In one embodiment of the present invention, processor 101 provides a flash copy operation that instantly copies all of the values from register file 103 into register file 104. This facilitates a rapid register checkpointing operation to support transactional execution.

[0052] Processor 101 also includes one or more functional units, such as adder 107 and multiplier 108. These functional units are used in performing computational operations involving operands retrieved from register files 103 or 104. As in a conventional processor, load and store operations pass through load buffer 111 and store buffer 112.
15

[0053] Processor 101 additionally includes a level one (L1) data cache 115, which stores data items that are likely to be used by processor 101. Note that each line in L1 data cache 115 includes a “load-marking bit,” which indicates that a data value from the line has been loaded during transactional execution. This load-marking bit is used to determine whether any interfering memory references take place during transactional execution as is described below with reference to FIGs. 3-8. Processor 101 also includes an L1 instruction cache (not shown).
20

[0054] Note that load-marking does not necessarily have to take place in L1 data cache 115. In general load-marking can take place at any level cache,
25

such as L2 cache 120. However, for performance reasons, the load-marking takes place at the cache level that is closest the processor as possible, which in this case is L1 data cache 115. Otherwise, loads would have to go to L2 cache 120 even on an L1 hit.

5 [0055] L2 cache 120 operates in concert with L1 data cache 115 (and a corresponding L1 instruction cache) in processor 101, and with L1 data cache 117 (and a corresponding L1 instruction cache) in processor 102. Note that L2 cache 120 is associated with a coherency mechanism 122, such as the reverse directory structure described in U.S. Patent Application No. 10/186,118, entitled, "Method and Apparatus for Facilitating Speculative Loads in a Multiprocessor System," filed on June 26, 2002, by inventors Shailender Chaudhry and Marc Tremblay (Publication No. US-2002-0199066-A1). This coherency mechanism 122 maintains "copyback information" 121 for each cache line. This copyback information 121 facilitates sending a cache line from L2 cache 120 to a requesting processor in cases where the current version of the cache line must first be retrieved from another processor.

10

15

[0056] Each line in L2 cache 120 includes a "store-marking bit," which indicates that a data value has been stored to the line during transactional execution. This store-marking bit is used to determine whether any interfering memory references take place during transactional execution as is described below with reference to FIGs. 3-8. Note that store-marking does not necessarily have to take place in L2 cache 120.

20

[0057] Ideally, the store-marking takes place in the cache level closest to the processor where cache lines are coherent. For write-through L1 data caches, writes are automatically propagated to L2 cache 120. However, if an L1 data cache is a write-back cache, we perform store-marking in the L1 data cache. (Note that the cache coherence protocol ensures that any other processor that

25

subsequently modifies the same cache line will retrieve the cache line from the L1 cache, and will hence become aware of the store-mark.)

Executing a Critical Section

5 [0058] FIG. 2 illustrates how a critical section is executed in accordance with an embodiment of the present invention. As is illustrated in the left-hand side of FIG. 2, a process that executes a critical section typically acquires a lock associated with the critical section before entering the critical section. If the lock has been acquired by another process, the process may have to wait until the other 10 process releases the lock. Upon leaving the critical section, the process releases the lock. (Note that the terms “thread” and “process” are used interchangeably throughout this specification.)

15 [0059] A lock can be associated with a shared data structure. For example, before accessing a shared data structure, a process can acquire a lock on the shared data structure. The process can then execute a critical section of code that accesses the shared data structure. After the process is finished accessing the shared data structure, the process releases the lock.

20 [0060] In contrast, in the present invention, the process does not acquire a lock, but instead executes a start transactional execution (STE) instruction before entering the critical section. If the critical section is successfully completed without interference from other processes, the process performs a commit operation, to commit changes made during transactional execution. This sequence of events is described in more detail below with reference to FIGs. 3-8.

25 [0061] Note that in one embodiment of the present invention a compiler replaces lock-acquiring instructions with STE instructions, and also replaces corresponding lock releasing instructions with commit instructions. (Note that there may not be a one-to-one correspondence between replaced instructions. For

example, a single lock acquisition operation comprised of multiple instructions may be replaced by a single STE instruction.) The above discussion presumes that the processor's instruction set has been augmented to include an STE instruction and a commit instruction. These instructions are described in more detail below with reference to FIGs. 3-9.

Transactional Execution Process

[0062] FIG. 3 presents a flow chart illustrating how transactional execution takes place in accordance with an embodiment of the present invention.

10 A process first executes an STE instruction prior to entering of a critical section of code (step 302). Next, the system transactionally executes code within the critical section, without committing results of the transactional execution (step 304).

15 [0063] During this transactional execution, the system continually monitors data references made by other processes, and determines if an interfering data access (or other type of failure) takes place during transactional execution. If not, the system atomically commits all changes made during transactional execution (step 308) and then resumes normal non-transactional execution of the program past the critical section (step 310).

20 [0064] On the other hand, if an interfering data access is detected, the system discards changes made during the transactional execution (step 312), and attempts to re-execute the critical section (step 314).

25 [0065] In one embodiment of the present invention, the system attempts the transactionally re-execute the critical section zero, one, two or more times. If these attempts are not successful, the system reverts back to the conventional technique of acquiring a lock on the critical section before entering the critical section, and then releasing the lock after leaving the critical section.

[0066] Note that an interfering data access can include a store by another process to a cache line that has been load-marked by the process. It can also include a load or a store by another process to a cache line that has been store-marked by the process.

- 5 **[0067]** Also note that circuitry to detect interfering data accesses can be easily implemented by making minor modifications to conventional cache coherence circuitry. This conventional cache coherence circuitry presently generates signals indicating whether a given cache line has been accessed by another processor. Hence, these signals can be used to determine whether an
10 interfering data access has taken place.

Starting Transactional Execution

[0068] FIG. 4 presents a flow chart illustrating a start transactional execution (STE) operation in accordance with an embodiment of the present invention. This flow chart illustrates what takes place during step 302 of the flow chart in FIG. 3. The system starts by checkpointing the register file (step 402). This can involve performing a flash copy operation from register file 103 to register file 104 (see FIG. 1). In addition to checkpointing register values, this flash copy can also checkpoint various state registers associated with the currently
15 executing process. In general, the flash copy operation checkpoints enough state to be able to restart the corresponding thread.
20

[0069] At the same time the register file is checkpointed, the STE operation also causes store buffer 112 to become “gated” (step 404). This allows existing entries in store buffer to propagate to the memory sub-system, but
25 prevents new store buffer entries generated during transactional execution from doing so.

[0070] The system then starts transactional execution (step 406), which involves load-marking and store-marking cache lines, if necessary, as well as monitoring data references in order to detect interfering references.

5 **Load-Marking Process**

[0071] FIG. 5 presents a flow chart illustrating how load-marking is performed during transactional execution in accordance with an embodiment of the present invention. During transactional execution of a critical section, the system performs a load operation. In performing this load operation if the load 10 operation has been identified as a load operation that needs to be load-marked, system first attempts to load a data item from L1 data cache 115 (step 502). If the load causes a cache hit, the system “load-marks” the corresponding cache line in L1 data cache 115 (step 506). This involves setting the load-marking bit for the cache line. Otherwise, if the load causes a cache miss, the system retrieves the 15 cache line from further levels of the memory hierarchy (step 508), and proceeds to step 506 to load-mark the cache line in L1 data cache 115.

Store-Marking Process

[0072] FIG. 6 presents a flow chart illustrating how store-marking is performed during transactional execution in accordance with an embodiment of the present invention. During transactional execution of a critical section, the system performs a store operation. If this store operation has been identified as a store operation that needs to be store-marked, the system first prefetches a corresponding cache line for exclusive use (step 602). Note that this prefetch 20 operation will do nothing if the line is already located in cache and is already in an exclusive use state. 25

[0073] Since in this example L1 data cache 115 is a write-through cache, the store operation propagates through L1 data cache 115 to L2 cache 120. The system then attempts to lock the cache line corresponding to the store operation in L2 data cache 115 (step 604). If the corresponding line is in L2 cache 120 (cache hit), the system “store-marks” the corresponding cache line in L2 cache 120 (step 610). This involves setting the store-marking bit for the cache line. Otherwise, if the corresponding line is not in L2 cache 120 (cache miss), the system retrieves the cache line from further levels of the memory hierarchy (step 608) and then proceeds to step 610 to store-mark the cache line in L2 cache 120.

[0074] Next, after the cache line is store-marked in step 610, the system enters the store data into an entry of the store buffer 112 (step 612). Note that this store data will remain in store buffer 112 until a subsequent commit operation takes place, or until changes made during the transactional execution are discarded.

[0075] Note that a cache line that is store marked by a given thread can be read by other threads. Note that this may cause the given thread to fail while the other threads continue.

20 Commit Operation

[0076] FIG. 7 presents a flow chart illustrating how a commit operation is performed after transactional execution completes successfully in accordance with an embodiment of the present invention. This flow chart illustrates what takes place during step 308 of the flow chart in FIG. 3.

[0077] The system starts by treating store-marked cache lines as though they are locked (step 702). This means other processes that request a store-

marked line must wait until the line is no longer locked before they can access the line. This is similar to how lines are locked in conventional caches.

[0078] Next, the system clears load-marks from L1 data cache 115 (step 704).

5 [0079] The system then commits entries from store buffer 112 for stores that are identified as needing to be marked, which were generated during the transactional execution, into the memory hierarchy (step 706). As each entry is committed, a corresponding line in L2 cache 120 is unlocked.

10 [0080] The system also commits register file changes (step 708). For example, this can involve functionally performing a flash copy between register file 103 and register file 104 in the system illustrated in FIG. 1.

Discarding Changes

[0081] FIG. 8 presents a flow chart illustrating how changes are discarded after transactional execution completes unsuccessfully in accordance with an embodiment of the present invention. This flow chart illustrates what takes place during step 312 of the flow chart in FIG. 3. The system first discards register file changes made during the transactional execution (step 802). This can involve either clearing or simply ignoring register file changes made during transactional execution. This is easy to accomplish because the old register values were checkpointed prior to commencing transactional execution. The system also clears load-marks from cache lines in L1 data cache 115 (step 804), and drains store buffer entries generated during transactional execution without committing them to the memory hierarchy (step 806). At the same time, the system unmarks corresponding L2 cache lines. Finally, in one embodiment of the present invention, the system branches to a target location specified by the STE

instruction (step 808). The code at this target location attempts to re-execute the critical section as is described above with reference to step 314 of FIG. 1.

Monitored Load Instructions

- 5 [0082] FIG. 9A presents a flow chart illustrating how monitored and unmonitored load instructions are generated in accordance with an embodiment of the present invention. This process takes place when a program is being generated to support transactional execution. For example, in one embodiment of the present invention, a compiler or virtual machine automatically generates 10 native code to support transactional execution. In another embodiment, a programmer manually generates code to support transactional execution.

- [0083] The system first determines whether a given load operation within a block of instructions to be transactionally executed needs to be monitored (step 902). In one embodiment of the present invention, the system determines 15 whether a load operation needs to be monitored by determining whether the load operation is directed to a heap. Note that a heap contains data that can potentially be accessed by other processes. Hence, loads from the heap need to be monitored to detect interference. In contrast, loads from outside the heap, (for example, from the local stack) are not directed to data that is shared by other processes, and 20 hence do not need to be monitored to detect interference.

- [0084] One embodiment of the present invention determines whether a load operation needs to be monitored at the programming-language level, by examining a data structure associated with the load operation to determine whether the data structure is a “protected” data structure for which loads need to 25 be monitored, or an “unprotected” data structure for which loads do not need to be monitored.

[0085] In yet another embodiment, the system allows a programmer to determine whether a load operation needs to be monitored.

[0086] If the system determines that a given load operation needs to be monitored, the system generates a “monitored load” instruction (step 904).

- 5 Otherwise, the system generates an “unmonitored load” instruction (step 906).

[0087] There are a number of different ways to differentiate a monitored load instruction from an unmonitored load instruction. (1) The system can use the op code to differentiate a monitored load instruction from an unmonitored load instruction. (2) Alternatively, the system can use the address of the load
10 instruction to differentiate between the two types of instructions. For example, loads directed to a certain range of addresses can be monitored load instructions, whereas loads directed to other address can be unmonitored load instructions.

[0088] Also note that an unmonitored load instruction can either indicate that no other process can possibly interfere with the load operation, or it can
15 indicate that interference is possible, but it is not a reason to fail. (Note that in some situations, interfering accesses to shared data can be tolerated.)

[0089] FIG. 9B presents a flow chart illustrating how monitored and unmonitored load instructions are executed in accordance with an embodiment of the present invention. The system first determines whether the load instruction is
20 a monitored load instruction or an unmonitored load instruction (step 910). This can be accomplished by looking at the op code of the load instruction, or alternatively, looking at the address for the load instruction. Note that the address can be examined by comparing the address against boundary registers, or possibly examining a translation lookaside buffer (TLB) entry fro the address to determine
25 if the address falls within a monitored range of addresses.

[0090] If the load instruction is a monitored load instruction, the system performs the corresponding load operation and load marks the associated cache

line (step 914). Otherwise, if the load instruction is an unmonitored load instruction, the system performs the load operation without load-marking the cache line (step 916).

- [0091] In another embodiment of the present invention, instead of
- 5 detecting interfering data accesses from other processes, the system does not allow a load operation from the current process cause other processes to fail. This can be accomplished by propagating additional information during the coherency transactions associated with the load operation to ensure that the load operation does not cause another process to fail.

10

Monitored Store Instructions

- [0092] FIG. 10A presents a flow chart illustrating how monitored and unmonitored store instructions are generated in accordance with an embodiment of the present invention. As was described above for load operations, this process
- 15 can take place when a compiler or virtual machine automatically generates native code to support transactional execution, or when a programmer manually generates code to support transactional execution.

- [0093] The system first determines whether a store operation within a block of instructions to be transactionally executed needs to be monitored
- 20 (step 1002). This determination can be made in the based on the same factors as for load instructions.

- [0094] If the system determines that a store operation needs to be monitored, the system generates a “monitored store” instruction (step 1004). Otherwise, the system generates an “unmonitored store” instruction (step 1006).

- 25 [0095] Note that monitored store instructions can be differentiated from unmonitored store instructions in the same way that monitored load instructions

21

can be differentiated from unmonitored load instructions, for example the system can use different op codes or different address ranges.

[0096] FIG. 10B presents a flow chart illustrating how monitored and unmonitored store instructions are executed in accordance with an embodiment of the present invention. The system first determines whether the store instruction is a monitored store instruction or an unmonitored store instruction (step 1010). This can be accomplished by looking at the op code for the store instruction, or alternatively, looking at the address for the store instruction. If the store instruction is a monitored store instruction, the system performs the corresponding store operation and store marks the associated cache line (step 1014). Otherwise, if the store instruction is an unmonitored store instruction, the system performs the store operation without store-marking the cache line (step 1016).

[0097] Note that a store-marked cache line can indicate one or more of the following: (1) loads from other processes to the cache line should be monitored; (2) stores from other processes to the cache line should be monitored; or (3) stores to the cache line should be buffered until the transactional execution completes.

[0098] In another embodiment of the present invention, instead of detecting interfering data accesses from other processes, the system does not allow a store operation from the current process cause another process to fail. This can be accomplished by propagating additional information during coherency transactions associated with the store operation to ensure that the store operation does not cause another process to fail.

[0099] The foregoing descriptions of embodiments of the present invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent

to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.